Usefulness of "Value of Statistical Life" in the Commercial Space Context: Limited, But Still Relevant

Contribution to "Anchoring and Black Swans:
Reconsidering Risk Aversion and the Future of Commercial Space",
NASA Contract #NX15AC03G – Document #6

Timothy Brennan*

Senior Fellow, Resources for the Future
Professor, School of Public Policy, University of Maryland-Baltimore County

brennan@rff.org

October 28, 2016

I. Introduction and context

The impetus for this project is the concern that public reactions to potential accidents, especially those that take lives—"black swans"—will be excessive. This concern is motivated by precedents regarding space-related accidents in the past (Apollo 1, Columbia, Challenger) and how people react to accidents in other contexts—"anchoring." As part of this project, James Bennett, co-principal investigator, provided a detailed examination of the role of anchors in forming public attitudes (Bennett 2016). Katrina McLaughlin (then) of Resources for the Future presented survey findings to a March 22, 2016 workshop of invited experts, convened by Molly Macauley, the initiator of principal investigator on this project.[1]

A major focus of the discussion[2] and this work is a tool public policy analysts employ for putting a monetary value on increased or reduced risk of death—the "value of statistical live" (VSL)—used to characterize when the benefit of safety regulation is worth the cost such regulation imposes. A more extensive description of VSL is below, but we can think of it as a number that encapsulates how much individuals are willing to pay—based on how

* This document in the package in a formal sense takes the place of the contribution that Molly Macauley, principal investigator, was unable to complete. It is therefore important to state the disclaimer that the views expressed here are not necessarily the views she would have expressed, or that she would have agreed with them. I benefited tremendously over the years from discussions with Dr. Macauley on this and many other issues involving how to think about the need for and design of public policies, but no one should infer her views from mine expressed here. It is everyone's loss that this component of this project is not the one Dr. Macauley would have contributed, reflecting the expertise, concern, thoughtfulness, and clarity that led her to create and design this project.

[1] McLaughlin's presentation, accompanying notes, and summary (without attribution) of workshop participants' thoughts on these findings are documents #2-4 in this package. Document #1 is the description of the workshop and list of expert participants.

[2] Document #4 provides notes from the discussion. References to the discussion throughout are to that document.

much they in fact pay—to reduce risks in their own daily activities. McLaughlin gave a presentation on VSL, including its differences among agencies and differences across countries, particularly differences correlated with average income of residents of those countries.[3]

This work was motivated by the contrast between the public attitudes toward space exploration—especially what Molly Macauley called "government risk conservatism"—and the use of a VSL measure to balance benefits and costs. This topic received some discussion by the participants in the March 22 workshop as well. My purpose here is to take a look at the limits of VSL and its potential relevance as we look toward increasing commercialization of space, in terms of both NASA use of commercial vehicles for its operations and fully private operations in space. We want to see what VSL and economic perspectives generally might say about future policies and regulations that involve commercial space risks, and which questions those perspectives do and do not answer.

After explaining the focus on mortality risk and defining and justifying VSL, the following discussion focuses on how two factors affect its relevance to commercial space exploration. The first is that those who take risks of commercial space programs are active participants, not part of a passive public that may be affected by policies that alter their mortality risk, e.g., by reducing emissions or putting guard rails on highways. The second is that the aversion to risk people express is not risk to themselves, but the risk of mortality to others. This leads to some thoughts on the relevance for setting policy in this area of possible error in these preferences and their potential malleability. I conclude with a couple of recommendations for policy and further research.

II. Focus on mortality risk

The focus of this document, and the project as a whole, is on risks to life that might accompany commercial space ventures. Commercial space ventures may present risks to property as well, such as a fire starting on the ground by launch-related material or problems presented by space debris (Macauley 1994, 2004). In principle, liability law can provide incentives to deter carelessness that could lead to the destruction of property, although statutory (rather than common law) assignments of liability for commercial launches are somewhat problematic (Brennan, Kousky and Macauley 2010).

Other factors also justify a focus on mortality risk. When monetized in VSL terms—the process described in the McLaughlin presentation and discussed further below—mortality risks can constitute the vast share of benefits from policies directed at protecting the public. For example, the Environmental Protection Agency (EPA), in a retrospective study of the benefits of the Clean Air Act (CAA) from 1970 to 1990 (EPA 1997) and a prospective study of CAA benefits from 1990 to 2020 (EPA 2011), found that 80-85% of the benefits come from reduced mortality risk. In addition, as the Bennett and McLaughlin contributions in this package discuss, it is death, and not property loss, that grabs the attention of

[3] See documents #2 and #3 in the package.

the public. Last and perhaps not least is the idea that some reductions in the risk of death are not worth the costs is, to many, a discomfiting if not reprehensible proposition.

III. The economists' measure – VSL

A. Background – benefit cost analysis principles

To understand how economists put a monetary value on reducing mortality risk—and thereby implying that reductions in that risk that cost more than this value are not worth doing—it is useful to step back and understand the economic approach to evaluating policy interventions in general.[4] VSL is the application of that approach to the particular case of policy interventions in settings where mortality risk is at stake. That commercial space enterprise may be one such setting is the rationale for this project.

In economics, the justification for a policy intervention is that the market fails to reach the efficient outcome. By "efficiency," we mean here the idea that if something is worth more to someone than it costs to produce it, that thing will be produced and that person will get it. There are a number of types of and causes for market failure. One is when competition is insufficient to prevent a seller to reap profits by charging price of a good or service above the cost of production, thereby preventing some buyers, willing to pay a lower price but still above cost, from getting the product.[5] A second failure, asymmetric information, is when one side of a market knows something about cost or quality that the other side does not, leading the ignorant side to assume the worst and causing the market to fall apart. The archetypal example is used cars, where the sellers know the quality of their cars but the buyers do not and cannot get that information, leading to only the lowest quality cars being sold (Akerlof 1970).

A third, externalities, is when hurdles such as unclear ownership rights, bargaining, search, and verification—known as transaction costs—prevent otherwise desirable exchanges from taking place. An example would be space debris. Some parties would pay more to reduce debris than it would cost space users to prevent debris, but those hurdles keep that transaction from taking place. This would lead to too much debris in critical orbits and justify policies to prevent or limit debris (Macauley 1994). Closely related to this concept are "public goods" that benefit many once produced but for which the provider either cannot charge (because it cannot exclude those who do not pay) or charges too much. Scientific research, including space exploration, is an example of a public good.

In all of these cases, the justification for policy intervention is that the benefits of some production or allocation of goods, services and resources—as measured by willingness to pay—exceed the costs of that production or allocation, but market failures preclude that production or allocation from taking place. Examples from space could include a lack of

[4] Again, the McLaughlin presentation and accompanying notes, documents #2 and #3 in this package, discuss VSL as well.

[5] There is a similar "mirror image" story when a buyer faces insufficient competition and sets a price too low.

competition in launch or satellite markets, users of launch vehicles or satellites being unable to verify vendor claims regarding quality, space debris, and damage from launches gone awry, or space exploration itself. On this account, the fundamental criterion for whether a policy intervention is good is whether the benefit of that intervention exceeds its cost. (In the absence of a market failure, all production or allocation where benefits exceed costs will have already taken place.)

The name for this test is benefit-cost analysis (BCA).[6] Executive Orders from Presidents Reagan, Clinton, and Obama have all endorsed the principle that regulations should be examined to see whether the benefits justify the costs.[7] If the willingness to pay for the benefits of an intervention, aggregated over all of the beneficiaries, exceeds the cost of the policy intervention, it should be undertaken. If the costs of the intervention exceed how much people would be willing to pay for the benefits, it is not worth doing. The beneficiaries would be better off with the money the intervention costs than with the intervention itself.

B. Application of BCA to mortality risk reduction – the value of statistical life

The starting point for evaluating the benefits of mortality risk reductions, or the costs of mortality risk increases, is to apply this question—how much are people willing to pay to reduce risk?[8] A leading method for assessing this risk begins with how much more people would need to be paid to accept a job with higher risks or, to put it another way, how much they would pay through reduced wages for a better paying job (Viscusi 1993). Other methods have included tradeoffs between earlier arrivals and speed-related highway risks, estimated costs of installing or using safety equipment (air bags, seat belts, fire detectors), and effects of air pollution on property values (Viscusi 1993 at 1936).

This kind of data is used in the following way. Suppose that labor market data indicate that people are willing to accept a reduced salary of $90 per month to avoid an additional risk of death of 1/100,000 during that month. If so, a group of 100,000 similarly situated people would pay $9 million in the aggregate ($90 times 100,000) to reduce an expected death in that group by one (1/100,000 times 100,000). This $9 million figure would be the

[6] In practice, doing BCA is not easy. By definition, there is some market failure, so one cannot look at prices directly to determine how much people or businesses would be willing to pay to eliminate air pollution, get better quality products, or for research conducted on the International Space Station. Controversies about how to weigh present costs against future benefits play a role in policies to address very long-term market failures, such as climate change brought on by greenhouse gas emissions (Goulder and Williams 2012). Measurement error and the value of preserving options can also complicate the decision as to whether the benefits of a policy exceed the costs (Farrow and Hayakawa 2002; Farrow 2012). Uncertainty and difficulty in quantifying the benefits or costs of some effects add further complexity (Farrow and Viscusi 2011).

[7] President Ronald Reagan, "Federal Regulation," Executive Order 12291, February 17, 1981; President William Jefferson Clinton, "Regulatory Planning and Review," Executive Order 12866, September 30, 1993; President Barack Obama, "Improving Regulation and Regulatory Review," Executive Order 13563, January 18, 2011. The Office of Management and Budget has set out for implementing benefit cost analysis in regulation; Office of Management and Budget, "Regulatory Analysis," Circular A-4, September 17, 2003.

[8] For purposes of this discussion I am avoiding controversies associated with potential discrepancies between how much people would pay to avoid a bad outcome and how much people would have to be paid to accept a good outcome. For more on this topic in the context of policy evaluation; see Brennan (2016).

aggregate willingness to pay to eliminate, one average, one death. It is this $9 million figure that is meant by the "value of a statistical life," or VSL. This "$90 dollar for 1/100,000 reduction in risk" example was not chosen by accident; the current VSL used by the Department of Transportation in estimating safety benefits is $9.1 million (Trottenberg and Rivkin 2013).

Before examining the relevance of a VSL number to commercial space risk, some important features and caveats need to be made. Perhaps foremost is that the VSL number is not based on the present value of lost earnings or replacement cost when someone dies. Rather, it is based specifically and only on willingness to pay to reduce risk. The two may be related, in that willingness to pay to avoid risk might be based to some degree how a person takes into account the value to his or her family of those earnings or non-monetary contributions, so higher earnings might be correlated with a higher VSL. But the VSL is not based on those lost earnings or replacement costs.

The example and use of a precise number provides an illusion of precision in VSL estimates. The determination of such a number, holding worker and other job characteristics constant, is rife with stochastic uncertainty. On top of that, extrapolating from a specific context, e.g., willingness to pay to get workers on some margin to forego income to take a less risky job, to the general population may be difficult if workers on that margin are unrepresentative, e.g., construction workers opting to build houses instead of skyscrapers. More broadly, and as Katrina McLaughlin's presentation illustrates, VSL is likely to vary across individuals based on income, age, other demographic characteristics, along with basic attitudes toward their own risk-taking.

For many, the most challenging aspect of the VSL idea is that it is a finite number, thereby implying a ceiling on how much should be spent on risk reduction. The derivation implies that if people are willing to pay only a finite amount to avoid risk—$90 in that example—the VSL will be a finite number. But this merely reflects experience People take on a small risk of a fatality for a benefit with a finite and not particularly large value, such as driving to see a movie or go out to dinner. Not everyone, even those with plenty of money, buys cars with the maximum conceivable amount of safety equipment. One may disagree with the $90 either as an estimate supported by the data or a reflection the people making choices understand at least roughly the effect on risk they are getting by giving up a higher wage. But the idea that we spend infinite amounts to avoid small increments in risk is routinely belied by observed behavior. Spending more than what those who benefit are willing to pay that through tax dollars or regulatory costs would provide a benefit with costs greater than that willingness to pay.

But while the conclusion that the VSL is finite seems reasonable, it seems to contradict with other observed behavior—that persons will pay as much as they can to save a particular life, including their own. However, there is no logical contradiction between a finite VSL and willingness to spend even more than that to avert a certain or highly probable death of oneself or those to whom one is close. The willingness to spend an unlimited amount to avoid a specific death can be thought of as the flip side of the coin that there is no compensation for death.

That inability to compensate is consistent with the observed behavior that people spend only a finite amount when the avoided risk is small. The VSL is not based on any notion of compensating for death or reflecting how much people would pay to avoid a certain or near certain death to themselves, family, or friends. It is just shorthand for the combined willingness to pay of a large set of people to avoid what to each of them is a small reduction in the risk of death. As such, it is fundamentally no different than how any other benefit is measured in a benefit cost analysis designed to correct for a market failure.

IV. But is VSL relevant?

The challenge before us arises because commercial space risk in some respects does not match the settings to which VSL or even benefit-cost analysis applies. Before introducing two important potential differences between commercial space risk and conventional policy settings, it may be useful to discuss one context in which VSL would be appropriate.

Suppose that the mortality risk associated with commercial space enterprise arises from the possibility that launch or re-entry could cause harm to bystanders on the ground. This "bystander risk" is just the type of setting in which VSL would apply. The market failure is that bystanders cannot readily contract with commercial space companies for compensation to bear the risk or to pay the space companies to take more care to protect the bystanders from injury.[9] Because of that, a regulator such as NASA or the Federal Aviation Administration might impose requirements to make launch and re-entry safer for bystanders.[10]

This bystander risk involves a small risk borne by third parties who would benefit from regulation to take the place of that to which they would have agreed, were such contracting possible. Such settings, as opposed to settings when specific lives are at stake with a high chance of harm, match those that generate the data that support VSL estimates. Those estimates in principle could be used to determine the point at which increasing the stringency of regulation would cost more than it would be worth to the protected parties.[11]

The discussion at the workshop in Document #4 pointed out two settings that lie outside the realm to which VSL measures are usually applied. The first involves risks to those

[9] That the market failure is symmetrical, that is, that the potentially harmed parties can neither require compensation from the potential injurer, nor pay the potential injurer to be more careful, is one of the insights of Coase (1960).

[10] In principle, we could also rely on *ex post* civil liability lawsuits, where space companies pay damages (based on VSL numbers) in the event of accidents, rather than *ex ante* regulation. The difficulty in litigation and the possibility that companies could avoid paying damages by going into bankruptcy may limit the effectiveness of liability lawsuits (Shavell 1986).

[11] The analysis in Documents #2 and #3 on variation of VSL across countries is relevant here. To the extent that bystanders at risk live in countries with considerably lower incomes than in the US, should the lower VSL number be used, thus reducing the appropriate amount space companies should put into making launches safer? The economic standard, not to provide benefits that cost more than what the beneficiaries would pay, suggest that a lower VSL would be appropriate. However, that hardly settles the ethical questions regarding different efforts to mitigate mortality risk based on income.

who choose to participate in commercial space enterprises—astronauts, workers, scientists, space tourists. When those bearing the risk already are contracting with those imposing the risk, they already may have a venue in which to pay for those risks that are worth avoiding. Consequently, there may not be a market failure warranting intervention, whether or not involving a VSL. This is where discussions regarding "informed consent" are particularly relevant.

A second feature of the commercial space context that differs from the usual VSL regulatory setting is that the relevant risk preference is the aversion of the public at large to deaths in space. This was a central feature of Katrina McLaughlin's presentations at the workshop, much of the discussion, and James Bennett's (2016) paper on anchoring that is Document #5 in this package. VSL and the data underlying it are designed to see how much people would pay to reduce risks to themselves. They do not at all speak to how much people are willing to pay to reduce risks to others, particularly strangers, as would generally be the case in commercial space ventures.

The next two sections discuss these challenges to the relevance of VSL. Both, particularly the latter, raise questions about the relevance of VSL in setting policy. I conclude with some suggestions for policy development and further research in light of these challenges.

V. Choosing the Risk: "Informed" Consent

The implication if not definition of a market failure is that one person or group of people, call it A, would like something, call it X, from another person or group of people, call it B, who could provide X to A at a lower cost than A is willing to pay but does not do so. In the context of commercial space risk, X would be the risk reduction, B would be the commercial space enterprise, and A would be the party or parties bearing the risk. The market failure is the inability to come up with a transaction in which B reduces risk to A, when the benefit of the risk reduction to A exceeds the added safety cost to B.

By this criterion, one reasonably doubt that there would be a market failure warranting policy concern in settings where A and B are already "transacting" with each other. In commercial space enterprise, this could entail a wider variety of potential "A" candidates. At early stages of development, one might have astronauts undertaking initial stages of testing. After verifying the viability of the flight itself, the enterprise might employ scientists and engineers for further assessment of the performance of the vehicles and equipment. As the commercial space sector matures, we would see employees to carry out much of the work, pilots, task specialists (e.g., low gravity miners), and other support workers. With space tourism, one would have passengers.

All of these parties have entered into employee or customer relationships with the commercial space service provider. These will involve more than a minimal description of the service provided to customers or the working conditions for the employees. They would likely include a list of amenities and qualities having to do with working or travel conditions. One might expect that among these amenities and qualities would be the level of safety. If that would work effectively, a reasonable conclusion would be that no external

regulation of commercial space enterprise safety is necessary. To use an economists' cliché, "the market would take care of it."

However, experience on the ground indicates that markets need not work very well. On the regulatory side, we see extensive regulation of the safety aspects of airplanes, automobiles, consumer appliances, electronic devices, access to pharmaceuticals, food content and handling, and workplace equipment and conditions. Tort law has developed doctrines on product liability and malpractice, in which the threat of having to pay damages for accidents leads manufacturers and service suppliers, ideally, to take the level of care that the potential victims would have taken had they the knowledge and control to ensure that the desired level of risk mitigation was provided.

It is that lack of knowledge and control that becomes relevant. To put the problem in terms expressed in the workshop, the "informed consent" of employees or customers regarding commercial space safety, necessary to justify leaving matters to agreements between these parties and the commercial space service provider, would likely be lacking. Most of these parties will lack the time, expertise, and equipment that it would take to determine and verify that the level of risk mitigation they would expect is the level of risk mitigation for which they would be willing to pay (or sacrifice earnings). It is one thing to imagine a Wright Brothers setting in which the risk takers are the same people who built the equipment with their own risk mind, and another to presume that the same knowledge and control would take place between commercial space service providers and their employees and customers. There may be consent, but it is unlikely to be informed.[12]

The problem created by inability to be reasonably informed is not just that employees and customers might be exposed to more risk than they would be if the commercial space service provider took into account their willingness to pay for reduced risk. Because mitigating risk is costly, it might be reasonable to expect that without ability to verify risk mitigation, these potential employees and customers might assume the worst—that commercial space service providers would provide no more safety than is warranted to protect the providers' own equipment. This could reduce demand for space tourism and lead employees to demand "worst case scenario" wage premiums, perhaps to the point of inhibiting the development of commercial space enterprise. The absence of policy may impede the commercial space sector more than excessively cautious policy.

As with commercial flight, safety regulation of commercial space enterprise may be warranted to provide the level of safety that employees and customers would be willing to buy for themselves. Regulators—a term I loosely use to include the FAA, DOT, NASA, and any other agency Congress might empower as commercial space enterprise grows—would require tools to decide when the benefits of safety regulations justify the cost of compliance, in terms of both direct costs and delay. It is generally in such contexts that VSL would

[12] Methods for estimating VSL from employment wage premiums or consumer demand for and use of safety equipment do presume that persons do know the risks they face, from which a willingness to pay for risk reduction can be inferred.

be used to monetize the benefits of reducing risks to parties that cannot mitigate that risk through direct transactions.

This, however, leaves open an important question: whose VSL counts? The examples in Bennett's (2016) historical study suggest that participants in many dangerous activities appear willing to take significant risks. Test pilots, from the early barnstorming days through the X-15 program, are notable but by no means unique examples. Mountain climbers, skydivers, race car drivers, and many others indicate a willingness to take risk that exceeds that of the average person.

One might expect the same of those who choose to participate in commercial space enterprise. It may be a mistake to conclude that these people have a lower VSL. It is not difficult to imagine that a test pilot or astronaut might show an average willingness to mitigate risk in normal contexts, such as driving a car with multiple air bags. Rather, the apparent tolerance for risk in the space contexts could reflect a willingness to accept more risk for the non-monetary benefits of being or working in space, whether contributing to science or personal adventure. Safety requirements based on a VSL, without taking the value of space participation, would be excessive and might likely remove opportunities for those few who would be going into space as part of commercial space enterprise.

The challenge then for NASA or other entities that might be charged with commercial space safety regulation to protect employees or customers can follow the procedures used to design safety regulations more generally. As such, VSL estimates can play a useful role. The contextual challenge for commercial space policy is not to come up with a different VSL for those likely to participate in commercial space ventures. Rather, it is to come up with ways to recognize and quantify the willingness of those participants to take on risk in order to contribute to science and to experience the wonders of space.

VI. Public dislike for participant risk

The above discussion showed that the VSL concept may be relevant in deciding how much regulation might be justified to mitigate risks borne by participants in commercial space enterprise, because of their own inability to become informed about the effect of the design of commercial spacecraft and systems that affect the risk they bear. This is analogous to regulations regarding consumer product, workplace, and food and drug safety. In the space setting, however, such regulation should be guided by the likelihood that participants in commercial space enterprise are more willing than members of the general public to accept higher risks in order to experience space and contribute to scientific knowledge and human advancement.

The general public becomes relevant when we turn from the participants' own attitudes toward risks to the attitudes of the general public regarding the risk those participants take. This is the central issue in the discussion of anchors for judging when risks become excessive. Examples include the public response to the fatal Challenger and Columbia shuttle accidents, particularly the former, on which teacher Christa McAuliffe was a passenger. This issue, described in detail in James Bennett's contribution to this project, was where

much of the experts' workshop focused. In the view of many of the participants, the driver of much of the policy intervention into flight safety and the expense and delay that followed has been driven not by participants but by pressure from the general public on Congress, the President, and directly or through them, on NASA.

Political forces on these institutions will be the manifestation of this pressure from the public. The ability of economics and policy norms generally to influence this process may be limited by a number of factors. First, and not surprisingly, economics expresses this pressure from the general public as a collective willingness to pay to mitigate risk. This can be quite large. If, say the average person in the US would pay $10 for a risk reduction, this becomes $3 billion over the entire population. Adding that amount into a benefit-cost analysis based solely on participant risk can lead to much more risk reduction than the industry and its participants would find appropriate. In that light, the concern of the experts is valid.

Because this affects the public at large, this collective willingness to pay will be based to some extent on the size of the population and its average income. The standard measure of the product of population and average income is the gross domestic product (GDP) statistic. Holding attitudes constant, this implies that as GDP grows, pressure from the general public to risk would grow as well. An attitude, encapsulated in the surveys Katrina McLaughlin presented, that the public accepts say a 1/100 disaster risk today might be a 1/200 risk as the country becomes larger and wealthier.

To the extent that willingness to pay of those outside the US becomes relevant to the discussion, the point of tolerable risk would fall even further. On the other hand, the other side of the coin is important. The benefits of space exploration are also appreciated by the general public. If the public's collective willingness to pay for those benefits increases in the same proportion as its collective willingness to pay to mitigate risks borne by commercial space participants, the level of risk acceptable to the public would not change.

Some normative criteria introduce questions about the extent to which general public aversion to space fatalities should play a role beyond that of promoting regulation reflecting the preferences of the participants themselves. One could construe the general public's aversion as being paternalistic. It does imply that a participant willing to take, say, a 1/1000 risk of death would be prevented from doing so if, taking the public's view into account, only missions with a 1/10,000 risk of death would be permitted. Thus, one could get a situation in which a commercial space enterprise has a willing participant, and where risk regulation is optimal from the participant perspective, but the enterprise is prohibited or delayed because the public finds it too risky. This is paternalistic, which in economics is among the harshest terms of opprobrium.[13]

[13] There is an analogous concept in philosophy, going back at least to Dworkin (1977 at 235-37). Dworkin made a distinction between self-regarding and other-regarding preferences, where an other-regarding preference is based on the well-being of someone else besides oneself. He argued that other-regarding preferences should not count in ethical assessments. If the other regarding preference is positive, e.g., X has a preference for Y's well-being, then Y's well-being gets double counted in ethical assessments, once for the direct effect on Y and second for its effect on X. He was also concerned with negative other-regarding preferences,

Workshop participants raised two factors suggesting that in the commercial space context, this quasi-paternalism may not be a problem. The first is that the public's willingness to get directly involved in commercial space operations, especially after accidents,[14] may be based on a lack of confidence in NASA or other regulatory controls over space risk. Many participants noted that the public might have more confidence if NASA would be more transparent regarding its risk policies and that it, and risk-assessment specialists, were more effective at communicating how benefits and risks are balanced.

This communication may be difficult. Not only is the concept of a VSL used to set that balance difficult on its own, but it runs up against a widespread aversion to quantifying anything having to with mortality. One can well imagine that a firm that used a VSL to determine that the benefits of a safety improvement were less than the cost would be found to have been cold and calculating and thus more likely to be held liable for a death because the safety improvement was not adopted. This could happen despite the economic principle that having businesses incorporate mortality risk into their decisions through the prospect of liability judgments would lead them to take the same level of care that their customers would have taken for themselves. Educational efforts not just on the likelihood of a risk—one participant observed that people exaggerate risk of a terrorist attack—but on how and even more, whether to quantify risk, will be no easy tasks.

The second and more promising reason that public distaste for participant risk may not be as large a factor as it has been in the context of NASA missions is the difference between risks taken by a public agency on behalf of the public and those undertaken by private interest for private gain. For NASA missions, astronauts are carrying out exploration and research missions on behalf of the public at large. The public supports these missions as taxpayers, and the knowledge and experienced gained are public goods. NASA's understandable efforts to engage the public in its missions likely promotes this sense of acting on behalf of the public. When there is a fatality, the public may bear a sense of responsibility to minimize the chance of future fatalities—even if NASA and its contractors have taken precautions that reduce risk to a level that astronauts were willing to accept.

As numerous participants in the expert panel stated, commercial space enterprise is quite different. It is not an activity undertaken by the public through its government. Rewards from the enterprise go to the providers of the space-related service and their customers. As one of the experts said in the workshop, the death of the test pilot on the Virgin Galactic received public attention but did not lead to an outcry that the government should delay space tourism. Another pointed out, with regard to space tourism, that the public may not get worked up about risks taken by "billionaires in space".

---

where X would prefer Y to be worse-off, because of envy or racism, for example, arguing that those should not justify reducing the weight given to Y's well-being.

[14] An interesting finding in Katrina McLaughlin's notes for her presentation is that some surveys showed that the public may be more willing to accept space-related risk after an accident. A possible explanation could be that some survey respondents might view demanding greater risk mitigation as, in essence, admitting a mistake in failing to insist upon more risk mitigation prior to the accident.

The public need not be dismissive altogether, in that if risks seem unreasonable then it may reach a conclusion that the regulators are not doing the job that they do in the many government contexts listed in Katrina McLaughlin's presentation and notes. If, many decades from now, space tourism and other activities become something available to the wider public, fatal accidents would receive the attention that fatal passenger airline or rail accidents do now—as something that could happen to anyone. However, this sort of intervention does not raise the prospect of "space exceptionalism" regarding accidents that were witnessed following the Apollo 1 and shuttle disaster. Perhaps the difference between NASA missions and commercial space enterprise will mean that the latter is not inappropriately hampered by public attitudes regarding risks to participants.

VII. Changing public risk preferences

A few times, the discussion in the experts' workshop touched on topics such as changing how the public regards risk. Following the title of the workshop and this project, this was usually put in terms of "anchoring". One way to put changing attitudes was "changing the anchor"; another expression was "unanchoring".

To the extent changing public attitudes follows from greater transparency and more effective communication regarding how the tradeoffs between risk and benefits are managed, it does not raise any fundamental conceptual problems. But changing public attitudes toward risk raises conceptual questions in two cases: when the public is not just misinformed but, by some standard, is wrong, and when the public's underlying preferences for risk, holding information constant, could be changed.

The idea that people make erroneous choices has been gaining traction in recent years, often under the name "behavioral economics". One of the experts in the workshop referred to pioneering work by Amos Tversky and Daniel Kahneman. Cass Sunstein, recently the head of the Office of Information and Regulatory Affairs and the chief overseer of regulation in the Executive Branch, and Richard Thaler, a professor in the University of Chicago Business School, published a bestseller, *Nudge*, about how to get people to make the choices they should have made rather than the choices they do make (Sunstein and Thaler 2009). Sunstein (2000) earlier suggested that one area in which the public is likely to be wrong is choices involving uncertainty.

Error, however, presents some difficulties for economic methods for balancing benefits against costs, including use of the VSL. Estimates of willingness to pay for benefits, such as reductions in mortality risk, are based on data people reveal through their choices. Behavioral economics claims, however, that the choices people make may not reveal their "true" preferences. The empirical foundation for benefit-cost analysis in general and in the case of morality risk in particular is undercut (Brennan 2014). Regarding the latter, if people make wrong choices regarding whether to take riskier jobs or purchase safety equipment, the data that inform regulatory evaluators of VSL are invalid.

If methods to change preferences are or were available, different but equally fundamental difficulties arise (Brennan 2006). If the basis for policy intervention is to create the out-

comes that would have transpired had market failures not prevented outcomes where the willingness to pay exceeds cost, should that be based on the willingness to pay before or after the preference change? If the former, why, if those preferences are no longer relevant? If the latter, could one justify any outcome where preferences are sufficiently malleable to support them? The prospect of changing preferences, as well as the prospect of error described above, makes unavoidable the issue of who gets to decide which preferences are choices are the right ones, if the public cannot be trusted to make the right decisions.

The difficulty of these issues suggests that absent very strong evidence to the contrary, it may be best to deal with the attitudes toward risk held by participants in commercial space enterprise and, to the extent it will play a role in commercial settings, the attitudes of the public toward fatal space-related accidents, rather than think about changing preferences apart from providing more information to the public to increase its confidence that regulators are adequately mitigating risk.

VIII. Summary and future directions

While controversial in assigning a monetary value to a reduction in the risk of mortality, and conceding difficulties in its estimation, the "value of statistical life" is a short-hand way to describe the application of the same benefit-cost analysis methods to the question of when the benefit of a regulation or policy is mortality risk. However, benefit-cost analysis and the "value of statistical life" are designed for settings when those who would benefit from a reduction in mortality risk have no say in how much risk they are willing to pay to mitigate because of a market failure. The commercial space context differs from these settings in two crucial ways. First, those whose risk could be reduced are already participants in the enterprise, not passive potential beneficiaries. Second, much of the concern about excessive regulation due to "space exceptionalism" arises because the policy driver is the public's interest in minimizing risk not to themselves but to participants in space ventures.

The first of these, that risks affects those who would be participating in the commercial space sector, need not obviate the need for regulation to ensure that commercial space enterprises mitigate risk. "Informed consent" may not be realistic to expect or infer. One need not expect those participants to have or be able to readily acquire the information to determine the appropriate level or risk or verify that appropriate safety measures have been taken. In principle, VSL is an appropriate tool to determine when the expected benefits of safety measures do and do not exceed the costs of those measures. However, the design of such regulations also should recognize the willingness of participants to take on risk because of their interest in experiencing space and contributing to scientific knowledge. A task for future research will be to come up with and incorporate estimates of that willingness to justify less stringent risk regulation than might be appropriate in other contexts, such as air travel or workplace safety.

The second, the interest of the general public in reducing risks to space travelers, is a more difficult issue. The attention paid by the experts participating on the March 22, 2016 workshop shows this. Because the general public can observe space fatalities, the public's potential aggregate can be almost unlimited. Dealing with the consequences of this prefer-

ence in and of itself will not be likely to be solved by more empirical work to come up with a better estimate.

Rather, it will likely come from one of two directions. One is that greater transparency and risk communication from NASA, other agencies, and the space community regulating space risk may help assure the public that regulations result from a reasonable balance of the benefits of commercial space enterprise and the risks to participants. This may assure members of the public that they need not feel like they have to pay more to mitigate risk beyond the level already set by the relevant regulators.

The second is that the general public's distaste for space fatality risk, particularly as revealed by pressure from their political representatives, is not likely to play as big a role in commercial space ventures. Some of the experts suggested this possibility during the workshop discussion. There is some reason to think that this may be the case. Commercial space sector participants are not acting on the public's behalf in the way that astronauts on government-funded space missions are carrying out the public's wishes. Fatalities when participants appear to have taken risks may garner some public attention but do not seem to lead to an outcry that can lead to higher costs and lengthy delays. However, only time and experience will tell.

References

Akerlof, George, "The Market for Lemons: Quality, Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 84 (1970): 488-500.

Bennett, James, "Historical Analogues as Anchoring References in Popular Perceptions of Spaceflight Risk" (2016), Document #5 in this package.

Brennan, Timothy, "Green Preferences as Regulatory Policy Instrument," *Ecological Economics*, vol. 56, no. 1 (2006): 144-54.

Brennan, Timothy, "Behavioral Economics and Policy Evaluation," *Journal of Benefit-Cost Analysis*, vol. 5, no. 1 (2014): 89-109.

Brennan, Timothy, "How Much Relevance Does Reality Imply? (Re)Considering the Endowment Effect**,** Resources for the Future Discussion Paper 16-31 (July 2016).

Brennan, Timothy, Carolyn Kousky and Molly Macauley, "Public-Private Co-Production of Risk: Government Indemnification of the Commercial Space Launch Industry", *Risk, Hazards & Crisis in Public Policy*, vol. 1, no. 1 (2010): 117-46.

Coase, Ronald, "The Problem of Social Cost", *Journal of Law and Economics* 3 (1960): 1-40.

Dworkin, Ronald, *Taking Rights Seriously* (Cambridge, MA: Harvard University Press, 1977).

Environmental Protection Agency, *The Benefits and Costs of the Clean Air Act, 1970 to 1990* (Washington: EPA, 1997).

Environmental Protection Agency, *The Benefits and Costs of the Clean Air Act from 1990 to 2020* (Washington: EPA, 2011).

Farrow, Scott, "A Missing Error Term in Benefit-Cost Analysis," *Environmental Science and Technology* 46 (2012): 2523-2528.

Farrow, Scott and Hiroshi Hayakawa, "Investing in Safety: An Analytical Precautionary Principle," *Journal of Safety Research* 33 (2002, August): 165-174.

Farrow, Scott and Kip Viscusi, ""Principles and Standards for the Benefit-Cost Analysis of Safety," *Journal of Benefit-Cost Analysis* 2 (2011) (no. 3, art. 5): 1-25.

Goulder, Lawrence and Roberton Williams III, "The Choice of Discount Rate for Climate Change Policy Evaluation," *Climate Change Economics* 03, (2012, 1250024) 1-18,

Macauley, Molly, "In Pursuit of a Sustainable Space Environment: Economic Issues in Regulating Space Debris," in Simpson, John (ed.), *Preservation of Near-Earth Space for Future Generations* (New York: Cambridge University Press, 1994): 147-58.

Macauley, Molly, "Economics of Space", in Sadeh, Eligar (ed.), *Space Politics and Policy, and Evolutionary Perspective* (Dordrecht: Springer, 2004): 181-200.

Shavell, Steven, "The Judgment Proof Problem," *International Review of Law and Economics* 6 (1986): 45-58.

Sunstein, Cass, "Cognition and Cost-Benefit Analysis," *The Journal of Legal Studies* 29 (2000): 1059-1103.

Sunstein, Cass and Richard Thaler, *Nudge: Improving Decisions about Health, Wealth, and Happiness* (London: Penguin Books, 2009).

Trottenberg, Polly and Robert Rivkin, "Guidance on Treatment of the Economic Value of a Statistical Life in U.S. Department of Transportation Analyses," Memo, U.S. Department of Transportation (2013), available at www.transportation.gov/sites/dot.dev/files/docs/VSL%20Guidance.doc.

Viscusi, Kip, "The Value of Risks to Life and Health," *Journal of Economic Literature* 31 (1993): 1912-46.